arXiv:1403.6501v1 [math.NT] 25 Mar 2014

COMPUTING POINTS OF BOUNDED HEIGHT IN PROJECTIVE SPACE OVER A NUMBER FIELD

DAVID KRUMM

ABSTRACT. We construct an algorithm for solving the following problem: given a number field K, a positive integer N, and a positive real number B, determine all points in $\mathbb{P}^N(K)$ having relative height at most B. A theoretical analysis of the efficiency of the algorithm is provided, as well as sample computations showing how the algorithm performs in practice. Two variants of the method are described, and examples are given to compare their running times. In the case $N = 1$ we compare our method to an earlier algorithm for enumerating elements of bounded height in number fields.

1. INTRODUCTION

Let K be a number field of degree n over $\mathbb{Q}$ with ring of integers $\mathcal{O}_K$, and let N be a positive integer. For any real number $B \geq 1$, define

$$\Omega(B) = \{P \in \mathbb{P}^N(K) : H_K(P) \leq B\},$$

where H_K is the relative height function on the set $\mathbb{P}^N(K)$. In [20] Schanuel proved that there is a constant c, depending only on classical invariants of K, such that $\#\Omega(B) \sim c \cdot B^{N+1}$ as $B \to \infty$. Thus, Schanuel's result provides a solution to the problem of estimating the number of points of bounded height in a projective space over K. In practice it can prove useful for various applications to have an algorithm corresponding to this counting problem, so that one can generate all points in $\Omega(B)$ for any given B. In the case $N = 1$, algorithms of this type have been used to compute bases for Mordell-Weil groups of elliptic curves [18]; to compute preperiodic points for quadratic polynomials [8]; and to find examples of abelian surfaces with everywhere good reduction over quadratic fields [7]. For larger values of N an algorithm does not seem to exist in the literature; the purpose of this article is to develop an efficient algorithm that can be used for any value of N.

A natural first approach to the problem of computing $\Omega(B)$ is to turn Schanuel's counting argument into an algorithm. Unfortunately, it is not clear that this solution can be implemented in practice. Schanuel reduces the counting problem to a question of estimating the number of lattice points inside a certain region in $\mathbb{R}^{n(N+1)}$ (see Schanuel's paper for details, or [15, Chap. 3, §5] for a sketch of the argument). While there are methods for computing lattice points inside bounded subsets of Euclidean space, these subsets must be relatively simple, and the dimension of the ambient space should be kept as small as possible. The region that occurs in Schanuel's paper is somewhat complicated (indeed, a substantial portion of the paper is spent on proving that the region is bounded and has a sufficiently smooth boundary), making it difficult to construct this region in a computer and determine all lattice points inside it.

The algorithm developed in the present article also involves a computation of lattice points in a bounded region; however, the ambient space can be taken to be either $\mathbb{R}^n$ or $\mathbb{R}^r$, where r is the rank of the unit group of $\mathcal{O}_K$, and the bounded region is in both cases a polytope. The cost of this simplification of the problem is not large: if we measure the efficiency of the method by comparing the size of $\Omega(B)$ to the size of the search space $\mathcal{S}(B)$ (that is, the set of all points generated by the algorithm while searching for points in $\Omega(B)$), then the inefficiency of the algorithm is bounded above by a constant as B tends to infinity.

Theorem 1.1. *There is a constant k, depending on N, K, and a choice of fundamental units in K, such that*

$$\limsup_{B\to\infty} \frac{\#\mathcal{S}(B)}{\#\Omega(B)} \leq k.$$

In the particular case $N = 1$, where other methods already exist, our asymptotic bounds on the size of the search space compare favorably to those of the other methods: the algorithm of Pethő and Schmitt [18] produces a search space that is larger than $\Omega(B)$ by a factor of B^{2n-2}; the algorithm given in [9] improves this to a factor of $(\log B)^r$. For the algorithm developed here, this factor is a constant independent of B.

In addition to the efficiency of our algorithm there is another salient feature to point out. Given a height bound B, the algorithm computes a subset $\mathcal{C}(B) \subset \mathcal{O}_K$ from which homogeneous coordinates of all points in $\Omega(B)$ can be taken. The set $\mathcal{C}(B)$ depends on B, K, and a choice of fundamental units in K, but is independent of the dimension of the ambient space $\mathbb{P}^N$. Thus, by taking tuples of elements of $\mathcal{C}(B)$ of appropriate length, one can determine all K-rational points of height bounded by B in any projective space over K.

This article is organized as follows. In §2 we set notation and record the theoretical results that are used in developing our algorithm. The main results of the paper are in §3, where we use the ideal class group of K to partition the set $\Omega(B)$ into subsets, and show how to reduce the computation of these subsets to a problem of finding lattice points in polytopes. In §4 we discuss ways of improving the efficiency of our method. The special case where K has only finitely many units is treated separately in §5. Asymptotics for the size of the search space of the algorithm are obtained in §6, and specific computations illustrating the performance of the algorithm are given in §7.

2. Background and notation

2.1. **Definition and computation of the relative height function.** Let M_K denote the set of nontrivial places of K. To every place $v \in M_K$ there corresponds an absolute value $|\ |_v$ on K extending one of the standard absolute values on $\mathbb{Q}$. We denote by K_v the completion of K with respect to $|\ |_v$, and by σ_v the natural inclusion $K \hookrightarrow K_v$. Restricting $|\ |_v$ to $\mathbb{Q}$ and completing we obtain a field $\mathbb{Q}_v$ which embeds into K_v yielding a finite extension $K_v/\mathbb{Q}_v$; the *local degree* of K at v is the degree n_v of this extension. The *normalized absolute value* on K corresponding to v is given by $\|x\|_v = |x|_v^{n_v}$. With this normalization we have the following product formula, which holds for every $x \in K^*$:

$$\prod_{v\in M_K} \|x\|_v = 1.$$

For a point $P = [x_0, \ldots, x_N] \in \mathbb{P}^N(K)$, the *height* of P (relative to K) is given by

$$H_K(P) = \prod_{v\in M_K} \max\{\|x_0\|_v, \ldots, \|x_N\|_v\}.$$

The product formula ensures that the height of P is independent of the choice of homogeneous coordinates for P. For purposes of explicit computations with absolute values and heights, a more concrete description of the places of K and of the function H_K will be needed.

Let M_K^∞ denote the set of places v for which $|\ |_v$ is Archimedean. For convenience, we will often write $v|\infty$ instead of $v \in M_K^\infty$. For every place $v|\infty$ the completion K_v can be identified with either $\mathbb{R}$ or $\mathbb{C}$, so that the inclusion $\sigma_v : K \hookrightarrow K_v$ is identified with either a real or complex embedding of K. The local degree n_v is 1 if $K_v = \mathbb{R}$ and 2 if $K_v = \mathbb{C}$. Hence, if the embeddings $K \hookrightarrow \mathbb{C}$ are explicitly known, one can use the relation $\|x\|_v = |\sigma_v(x)|_{\mathbb{C}}^{n_v}$ to compute, for any $x \in K$, all the numbers $\|x\|_v$ with $v|\infty$. Here, $|\ |_{\mathbb{C}}$ denotes the usual complex absolute value. This observation can be used to compute heights of points in $\mathbb{P}^N(K)$, as shown below.

We define a function $H_\infty : K^{N+1} \to \mathbb{R}$ by

$$H_\infty(x_0, \ldots, x_N) = \prod_{v|\infty} \max\{\|x_0\|_v, \ldots, \|x_N\|_v\}.$$

Let $\sigma_1, \ldots, \sigma_{r_1}$ be the real embeddings of K, and $\tau_1, \overline{\tau}_1, \ldots, \tau_{r_2}, \overline{\tau}_{r_2}$ the complex embeddings. It is a standard fact that the map $v \mapsto \sigma_v$ is a bijection between M_K^∞ and the set $\{\sigma_1, \ldots, \sigma_{r_1}, \tau_1, \ldots, \tau_{r_2}\}$. Using this fact we obtain the following alternate definition of H_∞ which is more suitable for computation:

$$H_\infty(x_0, \ldots, x_N) = \prod_{\sigma} \max\{|\sigma(x_0)|, \ldots, |\sigma(x_N)|\}, \tag{2.1}$$

where σ ranges over all embeddings $K \hookrightarrow \mathbb{C}$. It is well known that for any point $P = [x_0, \ldots, x_N] \in \mathbb{P}^N(K)$ we have the relation

$$H_K(P) = H_\infty(x_0, \ldots, x_N)/N(\mathfrak{a}), \tag{2.2}$$

where $\mathfrak{a}$ is the fractional ideal of $\mathcal{O}_K$ generated by $x_0, \ldots, x_N$. (For instance, see [21, p. 136, 3.7].) Here, $N(\mathfrak{a})$ denotes the norm of the ideal $\mathfrak{a}$. In practice, when computing $H_K(P)$ for a given point P, the formulas (2.2) and (2.1) will be used instead of our initial definition of the height.

2.2. **Results from Minkowski Theory.** Let $n = [K : \mathbb{Q}]$ be the degree of K over $\mathbb{Q}$. The *Minkowski embedding* of K is the map

$$\Phi : K \hookrightarrow \prod_{v|\infty} K_v \;\cong\; \mathbb{R}^n \tag{2.3}$$

given by $x \mapsto (\sigma_v(x))_v$. The stated isomorphism of real vector spaces is not canonical, so we will make a choice: we fix the ordering $\sigma_1, \ldots, \sigma_{r_1}, \tau_1, \ldots, \tau_{r_2}$ for the embeddings of K, thus obtaining an ordering $v_1, \ldots, v_{r_1+r_2}$ of the places $v|\infty$. This ordering induces an isomorphism

$$\prod_{v|\infty} K_v \;\cong\; \mathbb{R}^{r_1} \times \mathbb{C}^{r_2}.$$

Identifying $\mathbb{C}$ with $\mathbb{R}^2$ in the obvious way we obtain the isomorphism (2.3), and in particular a more concrete description of the Minkowski embedding:

$$\Phi(x) = (\sigma_1(x), \ldots, \sigma_{r_1}(x), \Re\tau_1(x), \Im\tau_1(x), \ldots, \Re\tau_{r_2}(x), \Im\tau_{r_2}(x)).$$

Under the map Φ, every nonzero ideal $\mathfrak{a}$ of $\mathcal{O}_K$ becomes a lattice of rank n in $\mathbb{R}^n$. More precisely, fixing an integral basis $\{\omega_1, \ldots, \omega_n\}$ for $\mathfrak{a}$, the vectors $\Phi(\omega_1), \ldots, \Phi(\omega_n)$ are linearly independent over $\mathbb{R}$ and generate $\Phi(\mathfrak{a})$ as a $\mathbb{Z}$-module. We denote by $F(\mathfrak{a})$ the fundamental parallelotope spanned by these vectors:

$$F(\mathfrak{a}) = \{c_1\Phi(\omega_1) + \cdots + c_n\Phi(\omega_n) : c_i \in [0, 1) \text{ for all } i\}. \tag{2.4}$$

The volume of $F(\mathfrak{a})$ is given by

$$\operatorname{Vol} F(\mathfrak{a}) = 2^{-r_2}|\Delta_K|^{1/2}N(\mathfrak{a}),$$

where Δ_K is the discriminant of K. Note that $F(\mathfrak{a})$ depends on the choice of integral basis for $\mathfrak{a}$, but its volume does not.

Let $r = r_1 + r_2 - 1$ be the rank of the unit group $\mathcal{O}_K^\times$. Recall the standard logarithmic map

$$\Lambda : K^* \longrightarrow \prod_{v|\infty} \mathbb{R} \;\cong\; \mathbb{R}^{r+1} \tag{2.5}$$

given by $x \mapsto (\log \|x\|_v)_v$. Again, the above isomorphism is not canonical, so we choose the isomorphism induced by the ordering $\sigma_1, \ldots, \sigma_{r_1}, \tau_1, \ldots, \tau_{r_2}$. Thus, we obtain the more concrete description

$$\Lambda(x) = (\log|\sigma_1(x)|, \ldots, \log|\sigma_{r_1}(x)|, 2\log|\tau_1(x)|, \ldots, 2\log|\tau_{r_2}(x)|).$$

A classical theorem of Dirichlet states that the image of $\mathcal{O}_K^\times$ under the map Λ is a lattice of rank r in the hyperplane consisting of all points (t_v) such that $\sum_v t_v = 0$. Any collection of units $\boldsymbol{\varepsilon} = \{\varepsilon_1, \dots, \varepsilon_r\}$ such that the vectors $\Lambda(\varepsilon_j)$ form a basis for the lattice $\Lambda(\mathcal{O}_K^\times)$ is called a *system of fundamental units* in K. We fix a choice of fundamental units and define

$$F(\boldsymbol{\varepsilon}) = \{t_1\Lambda(\varepsilon_1) + \cdots + t_r\Lambda(\varepsilon_r) : |t_j| \le 1/2 \text{ for all } j\}. \tag{2.6}$$

Note that $F(\boldsymbol{\varepsilon})$ is the closure of a fundamental domain for the lattice $\Lambda(\mathcal{O}_K^\times)$. In each direction v we will need to consider how far from the origin the vectors in $F(\boldsymbol{\varepsilon})$ can be; thus, we define numbers D_v by

$$D_v = \max_{\eta \in F(\boldsymbol{\varepsilon})} \eta_v. \tag{2.7}$$

We will also denote D_v by D_σ if v corresponds to the embedding $\sigma \in \{\sigma_1, \dots, \sigma_{r_1}, \tau_1, \dots, \tau_{r_2}\}$. Using the finite subset

$$V(\boldsymbol{\varepsilon}) = \{t_1\Lambda(\varepsilon_1) + \cdots + t_r\Lambda(\varepsilon_r) : |t_j| = 1/2 \text{ for all } j\}$$

we obtain a different description of the numbers D_v that is better for computation:

$$D_v = \max_{\eta \in V(\boldsymbol{\varepsilon})} \eta_v. \tag{2.8}$$

3. Points of bounded height with fixed ideal class

The starting point for our method is the observation that the set $\mathbb{P}^N(K)$ can be divided into ideal classes: to every point $P = [x_0, \dots, x_N] \in \mathbb{P}^N(K)$ there corresponds the ideal class $\mathrm{Cl}(P)$ of the fractional ideal generated by $x_0, \dots, x_N$; this is independent of the choice of homogeneous coordinates for P. Since there are only finitely many ideal classes of $\mathcal{O}_K$, this observation reduces the problem of finding all points in $\Omega(B)$ to the following:

Given a nonzero ideal $\mathfrak{a}$ of $\mathcal{O}_K$, compute the set

$$\Omega(\mathfrak{a}, B) = \{P \in \mathbb{P}^N(K) : \mathrm{Cl}(P) = \mathrm{Cl}(\mathfrak{a}) \text{ and } H_K(P) \le B\}.$$

We will therefore begin by considering this specialized version of the main problem. We assume here that the unit rank r is positive; the simpler case when $r = 0$ is treated in §5.

3.1. **A search space for $\Omega(\mathfrak{a}, B)$.** The main result of this section, namely Theorem 3.1 below, will allow us to construct a finite set of points containing $\Omega(\mathfrak{a}, B)$. Once this larger set is known, one can eliminate extraneous points from it by computing their heights and comparing to the bound B.

We define a region

$$\mathcal{P}(\mathfrak{a}, B) \subset \prod_{v|\infty} K_v \;\cong\; \mathbb{R}^n$$

as follows: letting D_v be the number (2.7) for every $v|\infty$, the set $\mathcal{P}(\mathfrak{a}, B)$ consists of all points $(s_v) \in \prod_{v|\infty} K_v$ such that

$$|s_v| \le (B \cdot N(\mathfrak{a}))^{1/n} \exp(D_v/n_v) \qquad \forall\, v|\infty.$$

Note that as a subset of $\mathbb{R}^n$, $\mathcal{P}(\mathfrak{a}, B)$ is a cartesian product of r_1 closed intervals in $\mathbb{R}$ and r_2 closed disks in $\mathbb{R}^2$. More precisely, $\mathcal{P}(\mathfrak{a}, B)$ consists of the points

$$(a_1, \dots, a_{r_1}; x_1, y_1, \dots, x_{r_2}, y_{r_2}) \in \mathbb{R}^n$$

such that

$$|a_i| \le (B \cdot N(\mathfrak{a}))^{1/n} \exp(D_{\sigma_i}) \;\text{ and }\; x_j^2 + y_j^2 \le (B \cdot N(\mathfrak{a}))^{2/n} \exp(D_{\tau_j})$$

for all indices $1 \le i \le r_1$ and $1 \le j \le r_2$.

Theorem 3.1. *For every point $P \in \Omega(\mathfrak{a}, B)$ there exist $x_0, \ldots, x_N \in \mathfrak{a}$ such that the following hold:*

- $P = [x_0, \ldots, x_N]$;
- $\mathfrak{a} = (x_0, \ldots, x_N)$;
- $|N_{K/\mathbb{Q}}(x_i)| \leq B \cdot N(\mathfrak{a})$ *for all i;*
- $\Phi(x_i) \in \mathcal{P}(\mathfrak{a}, B)$ *for all i.*

In order to prove the theorem we will need some auxiliary results.

Lemma 3.2. *Define a map $\theta : K^{N+1} \setminus 0 \longrightarrow \prod_{v|\infty} \mathbb{R}$ by*

$$\theta(\vec{\alpha}) = \left(\log \max_i \|\alpha_i\|_v \right)_v,$$

where $\vec{\alpha} = (\alpha_0, \ldots, \alpha_N)$. Then θ has the following properties:

- *For every $\vec{\alpha} \in K^{N+1} \setminus 0$, $\sum_v \theta(\vec{\alpha})_v = \log H_\infty(\vec{\alpha})$.*
- *If u is a unit in $\mathcal{O}_K$, then $\theta(u\vec{\alpha}) = \Lambda(u) + \theta(\vec{\alpha})$.*

Proof. Both properties follow immediately from the definitions. □

Lemma 3.3. *For any $\lambda \in K$ and $Y \in K^{N+1}$ we have $H_\infty(\lambda Y) = |N_{K/\mathbb{Q}}(\lambda)| \cdot H_\infty(Y)$.*

Proof. This is a consequence of (2.1). □

Proof of Theorem 3.1. Let $P \in \Omega(\mathfrak{a}, B)$. Since $\mathrm{Cl}(P) = \mathrm{Cl}(\mathfrak{a})$, there are homogeneous coordinates $[y_0, \ldots, y_N]$ for P such that $\mathfrak{a} = (y_0, \ldots, y_N)$. Letting $Y = (y_0, \ldots, y_N) \in K^{N+1}$, we have

$$H_\infty(Y) = H_K(P) \cdot N(\mathfrak{a}) \leq B \cdot N(\mathfrak{a}).$$

The vectors $\Lambda(\varepsilon_1), \ldots, \Lambda(\varepsilon_r)$ together with the vector (n_v) of local degrees form a basis for the Euclidean space $\prod_{v|\infty} \mathbb{R}$, so we can write

$$\theta(Y) = t \cdot (n_v) + \sum_{j=1}^{r} t_j \cdot \Lambda(\varepsilon_j) \tag{3.1}$$

for some real numbers $t, t_1, \ldots, t_r$. Considering the sum of the coordinates of the vectors on both sides of this equation, we find that

$$t = \log H_\infty(Y)^{1/n}.$$

Here, we are using the first property of the map θ listed in Lemma 3.2. Let n_j be an integer closest to t_j for every j, so that $|t_j - n_j| \leq 1/2$, and let $u = \varepsilon_1^{-n_1} \cdots \varepsilon_r^{-n_r} \in \mathcal{O}_K^\times$. We now choose a different set of homogeneous coordinates for P: set $x_i = uy_i$ and $X = (x_0, \ldots, x_N) \in K^{N+1}$. Clearly, $P = [x_0, \ldots, x_N]$ and $\mathfrak{a} = (x_0, \ldots, x_N)$; moreover, using Lemma 3.3 we see that $H_\infty(X) = H_\infty(uY) = H_\infty(Y)$, so our work above shows that

$$H_\infty(X) \leq B \cdot N(\mathfrak{a}) \quad \text{and} \quad t = \log H_\infty(X)^{1/n}.$$

It follows from (2.1) that for every index i we have $|N_{K/\mathbb{Q}}(x_i)| \leq H_\infty(X) \leq B \cdot N(\mathfrak{a})$. It remains to show that $\Phi(x_i) \in \mathcal{P}(\mathfrak{a}, B)$ for all i.

Using the second property of θ stated in Lemma 3.2 we obtain by (3.1) that

$$\theta(X) = \theta(uY) = \Lambda(u) + \theta(Y) = t \cdot (n_v) + \eta$$

for some $\eta \in F(\boldsymbol{\varepsilon})$. Considering the equation $\theta(X) = t \cdot (n_v) + \eta$ one coordinate at a time we find that

$$\log \max_i \|x_i\|_v = \log H_\infty(X)^{n_v/n} + \eta_v \qquad \forall\, v|\infty,$$

and so

$$\max_i \|x_i\|_v = H_\infty(X)^{n_v/n} \exp(\eta_v) \quad \forall\, v|\infty.$$

Thus, for every index $i \in \{0, \dots, N\}$ we have

$$\|x_i\|_v \le H_\infty(X)^{n_v/n} \exp(\eta_v) \le (B \cdot N(\mathfrak{a}))^{n_v/n} \exp(D_v) \quad \forall\, v|\infty.$$

By definition, this means that $\Phi(x_i) \in \mathcal{P}(\mathfrak{a}, B)$ for every i. □

Define a subset $\mathcal{C}(\mathfrak{a}, B) \subset \mathcal{O}_K$ by

$$\mathcal{C}(\mathfrak{a}, B) = \{x \in \mathfrak{a} : \Phi(x) \in \mathcal{P}(\mathfrak{a}, B) \text{ and } |N_{K/\mathbb{Q}}(x)| \le B \cdot N(\mathfrak{a})\}.$$

Note that $\mathcal{C}(\mathfrak{a}, B)$ is finite, since there can only be finitely many points of the lattice $\Phi(\mathfrak{a})$ lying inside the bounded region $\mathcal{P}(\mathfrak{a}, B)$. Theorem 3.1 shows that every point $P \in \Omega(\mathfrak{a}, B)$ has homogeneous coordinates coming from $\mathcal{C}(\mathfrak{a}, B)$, and thus provides a finite search space for the points in $\Omega(\mathfrak{a}, B)$. More precisely, we have the following description of $\Omega(\mathfrak{a}, B)$.

Corollary 3.4. *The set $\Omega(\mathfrak{a}, B)$ consists of all points of the form $P = [x_0, \dots, x_N]$ satisfying*

- *$x_i \in \mathcal{C}(\mathfrak{a}, B)$ for all i;*
- *$\mathfrak{a} = (x_0, \dots, x_N)$; and*
- *$H_\infty(x_0, \dots, x_N) \le B \cdot N(\mathfrak{a})$.*

Proof. Let $P \in \Omega(\mathfrak{a}, B)$ be any point. By Theorem 3.1 we know that there exist elements $x_0, \dots, x_N \in \mathcal{C}(\mathfrak{a}, B)$ such that $P = [x_0, \dots, x_N]$ and $\mathfrak{a} = (x_0, \dots, x_N)$. Moreover, since $H_K(P) \le B$, then (2.2) implies that $H_\infty(x_0, \dots, x_N) \le B \cdot N(\mathfrak{a})$.

Conversely, suppose that $P = [x_0, \dots, x_N]$ is any point whose coordinates satisfy the conditions listed in the corollary. The last two conditions then imply that $\mathrm{Cl}(P) = \mathrm{Cl}(\mathfrak{a})$ and $H_K(P) \le B$, so that $P \in \Omega(\mathfrak{a}, B)$ by definition. □

Using this corollary we obtain the following initial step towards our main algorithm.

Algorithm 1. Computing $\Omega(\mathfrak{a}, B)$.

(1) Create an empty list L.
(2) Compute the set $\mathcal{C}(\mathfrak{a}, B)$.
(3) For every tuple $(x_0, \dots, x_N)$ of elements of $\mathcal{C}(\mathfrak{a}, B)$:
    If $\mathfrak{a} = (x_0, \dots, x_N)$, then:
    (a) Let $P = [x_0, \dots, x_N] \in \mathbb{P}^N(K)$.
    (b) If $H_\infty(x_0, \dots, x_N) \le B \cdot N(\mathfrak{a})$, then include P in L.
(4) Return L.

3.2. **Exploiting group actions.** Before considering how the steps of Algorithm 1 can be carried out in practice, we discuss a modification of the crucial step in the algorithm, which is to compute the set $\mathcal{C}(\mathfrak{a}, B)$. Once this set has been computed, the next step in Algorithm 1 is to build $(N+1)$-tuples of elements of $\mathcal{C}(\mathfrak{a}, B)$ and check the heights of the corresponding points in $\mathbb{P}^N(K)$. We will show here that it is possible to replace $\mathcal{C}(\mathfrak{a}, B)$ by a proper subset; this will have the effect of reducing significantly the number of tuples that need to be considered when computing $\Omega(\mathfrak{a}, B)$. In addition, we show how to use two group actions to reduce the number of height computations that are carried out.

Let μ_K denote the group of roots of unity in K. From the definitions it follows that if $x \in \mathcal{C}(\mathfrak{a}, B)$ and $\zeta \in \mu_K$, then $\zeta x \in \mathcal{C}(\mathfrak{a}, B)$; hence, the group μ_K acts on $\mathcal{C}(\mathfrak{a}, B)$. Let $g_0, g_1, \dots, g_t$ be elements representing all the orbits of this action. We will see that in Algorithm 1, rather than considering all $(N+1)$-tuples of elements

of $\mathcal{C}(\mathfrak{a}, B)$, one can restrict attention to tuples of elements of the set $\{g_0, g_1, \dots, g_t\}$. Further reductions can be achieved by using two group actions on the set

$$K_\bullet^{N+1} = K^{N+1} \backslash \{(0, \dots, 0)\}.$$

First, the action of the symmetric group S_{N+1}: for $\pi \in S_{N+1}$ and $X = (x_0, \dots, x_N) \in K_\bullet^{N+1}$,

$$\pi \cdot X = (x_{\pi^{-1}(0)}, \dots, x_{\pi^{-1}(N)}).$$

Second, the action of the group μ_K^N: for $u = (\zeta_0, \dots, \zeta_{N-1}) \in \mu_K^N$ and $X = (x_0, \dots, x_N) \in K_\bullet^{N+1}$,

$$u \cdot X = (\zeta_0 x_0, \dots, \zeta_{N-1} x_{N-1}, x_N).$$

For $X \in K_\bullet^{N+1}$ we let

$$\mathcal{O}(X) = \{[u \cdot (\pi \cdot X)] : u \in \mu_K^N \text{ and } \pi \in S_{N+1}\} \subset \mathbb{P}^N(K). \tag{3.2}$$

Here, we denote by $[Y]$ the equivalence class in $\mathbb{P}^N(K)$ of the point $Y \in K_\bullet^{N+1}$.

Proposition 3.5. *Let $g_0, g_1, \dots, g_t \in \mathcal{C}(\mathfrak{a}, B)$ be elements representing all the orbits of the action of μ_K on $\mathcal{C}(\mathfrak{a}, B)$. Let $\mathcal{M}$ be the set of all tuples $X = (g_{i_0}, \dots, g_{i_N})$ such that:*

- $0 \le i_0 \le i_1 \le \cdots \le i_N \le t$;
- $\mathfrak{a} = (g_{i_0}, \dots, g_{i_N})$; *and*
- $H_\infty(X) \le B \cdot N(\mathfrak{a})$.

Then

$$\Omega(\mathfrak{a}, B) = \bigcup_{X \in \mathcal{M}} \mathcal{O}(X).$$

Proof. The definition of $\mathcal{M}$ implies that for any $X \in \mathcal{M}$, the point $P = [X]$ satisfies $H_K(P) \le B$ and $\mathrm{Cl}(P) = \mathrm{Cl}(\mathfrak{a})$, so that $P \in \Omega(\mathfrak{a}, B)$. A simple calculation shows that all the points in $\mathcal{O}(X)$ have the same height as P and the same ideal class; therefore, $\mathcal{O}(X) \subseteq \Omega(\mathfrak{a}, B)$. This proves one inclusion in the proposition.

To see the reverse inclusion, let $P \in \Omega(\mathfrak{a}, B)$. By Theorem 3.1 we can write $P = [y_0, \dots, y_N]$ with $y_i \in \mathcal{C}(\mathfrak{a}, B)$ and $\mathfrak{a} = (y_0, \dots, y_N)$. For every index i we have $y_i = z_i g_{e_i}$ for some $z_i \in \mu_K$ and some index $e_i \in \{0, \dots, t\}$; hence,

$$P = [z_0 g_{e_0}, \dots, z_N g_{e_N}].$$

Letting $\zeta_i = z_i z_N^{-1}$ for $i \in \{0, \dots, N-1\}$ we obtain

$$P = [\zeta_0 g_{e_0}, \dots, \zeta_{N-1} g_{e_{N-1}}, g_{e_N}].$$

By applying a permutation π of the set $\{0, \dots, N\}$ we may arrange the indices e_i so that

$$e_{\pi(0)} \le \cdots \le e_{\pi(N)}.$$

Define elements $x_0, \dots, x_N$ by $x_i = g_{e_{\pi(i)}}$, so that $x_{\pi^{-1}(j)} = g_{e_j}$. Then

$$P = [\zeta_0 x_{\pi^{-1}(0)}, \dots, \zeta_{N-1} x_{\pi^{-1}(N-1)}, x_{\pi^{-1}(N)}].$$

Letting $X = (x_0, \dots, x_N) \in K_\bullet^{N+1}$ and $u = (\zeta_0, \dots, \zeta_{N-1}) \in \mu_K^N$, this means by definition that

$$P = [u \cdot (\pi \cdot X)] \in \mathcal{O}(X).$$

We claim that $X \in \mathcal{M}$. First of all, we have $X = (g_{e_{\pi(0)}}, \dots, g_{e_{\pi(N)}})$ with increasing indices $e_{\pi(i)}$. Furthermore, from the construction of X it follows that the entries $x_0, \dots, x_N$ are associate – in some order – to

$y_0, \dots, y_N$. In particular, the ideal generated by $x_0, \dots, x_N$ is equal to the ideal $(y_0, \dots, y_N)$, which is $\mathfrak{a}$ by assumption. Finally, since all the points in $\mathcal{O}(X)$ have the same height, then

$$H_\infty(X)/N(\mathfrak{a}) = H_K([X]) = H_K(P) \leq B,$$

and hence $H_\infty(X) \leq B \cdot N(\mathfrak{a})$. This shows that $X \in \mathcal{M}$, which completes the proof of the proposition. □

We deduce from Proposition 3.5 the following improvement of Algorithm 1 in which fewer points need to be considered and fewer height computations are carried out.

Algorithm 2. Computing $\Omega(\mathfrak{a}, B)$.

(1) Create an empty list L.
(2) Compute representatives $g_0, g_1, \dots, g_t$ of the orbit space $\mathcal{C}(\mathfrak{a}, B)/\mu_K$.
(3) For every tuple of indices $(i_0, \dots, i_N)$ such that $0 \leq i_0 \leq i_1 \leq \cdots \leq i_N \leq t$:
    (a) Let X be the point $(g_{i_0}, \dots, g_{i_N})$.
    (b) If $\mathfrak{a} = (g_{i_0}, \dots, g_{i_N})$ and $H_\infty(X) \leq B \cdot N(\mathfrak{a})$, then:
    Include in L all the points in $\mathcal{O}(X)$.
(4) Return L.

Remark 3.6. With the notation of Proposition 3.5, it may very well happen that the sets $\mathcal{O}(X)$ and $\mathcal{O}(Y)$ for distinct tuples $X, Y \in \mathcal{M}$ have a nonempty intersection. Furthermore, it may happen that the points of the form $[u \cdot (\pi \cdot X)]$ constructed when computing $\mathcal{O}(X)$ are not all distinct. Hence, in Algorithm 2 one may want to remove duplicate points from the list L before returning it.

Remark 3.7. Everything stated in this section remains valid if the group μ_K is replaced throughout by its subgroup $\{\pm 1\}$. This observation will be used below to give two different versions of our main algorithm.

3.3. **Computing $\mathcal{C}(\mathfrak{a}, B)$.** We turn now to the question of how the set $\mathcal{C}(\mathfrak{a}, B)$ can be computed. Two approaches to this problem will be discussed, and in each case it will be shown that one can compute either $\mathcal{C}(\mathfrak{a}, B)/\mu_K$ or $\mathcal{C}(\mathfrak{a}, B)/\{\pm 1\}$ instead of the entire set $\mathcal{C}(\mathfrak{a}, B)$. By Proposition 3.5 and Remark 3.7, having a complete set of representatives for either one of these orbit spaces is enough to determine $\Omega(\mathfrak{a}, B)$ using Algorithm 2.

3.3.1. *First approach.* The elements of $\mathcal{C}(\mathfrak{a}, B)$ can be found by first computing the larger set

$$\widetilde{\mathcal{C}}(\mathfrak{a}, B) = \{x \in \mathfrak{a} : \Phi(x) \in \mathcal{P}(\mathfrak{a}, B)\}$$

and then eliminating those elements of $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ whose norms exceed the bound $B \cdot N(\mathfrak{a})$. Note that the computation of $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ is equivalent to finding all points of the lattice $\Phi(\mathfrak{a})$ that lie inside the region $\mathcal{P}(\mathfrak{a}, B)$; it would therefore be desirable to make this region as small as possible. The size of $\mathcal{P}(\mathfrak{a}, B)$ is tied to the sizes of the numbers D_v for $v|\infty$, which in turn are determined by the lengths of the vectors $\Lambda(\varepsilon_1), \dots, \Lambda(\varepsilon_r)$ forming a basis for the lattice $\Lambda(\mathcal{O}_K^\times) \subset \mathbb{R}^{r+1}$. Thus, a basis consisting of short vectors should be computed. This can be achieved by first computing any system $\{\varepsilon_1, \dots, \varepsilon_r\}$ of fundamental units (for instance, using the method described in [5, §6.5.3]) and then applying the LLL reduction algorithm [16], or faster variants such as the Nguyen-Stehlé algorithm [17], to the basis $\Lambda(\varepsilon_1), \dots, \Lambda(\varepsilon_r)$. Having obtained a set of "short" fundamental units in this way, the corresponding numbers D_v can be computed using (2.8). The region $\mathcal{P}(\mathfrak{a}, B)$ is then determined, and we need to find all points of the lattice $\Phi(\mathfrak{a})$ that lie inside it. For this purpose it will be convenient to work with a slightly larger region defined as follows. Recall that $\mathcal{P}(\mathfrak{a}, B)$ consists of the points

$$(a_1, \dots, a_{r_1}; x_1, y_1, \dots, x_{r_2}, y_{r_2}) \in \mathbb{R}^n$$

such that

$$|a_i| \leq (B \cdot N(\mathfrak{a}))^{1/n} \exp(D_{\sigma_i}) \text{ and } x_j^2 + y_j^2 \leq (B \cdot N(\mathfrak{a}))^{2/n} \exp(D_{\tau_j})$$

for all indices $1 \leq i \leq r_1$ and $1 \leq j \leq r_2$. We define $\mathcal{P}'(\mathfrak{a}, B) \subset \mathbb{R}^n$ to be the region consisting of all points

$$(a_1, \ldots, a_{r_1}; x_1, y_1, \ldots, x_{r_2}, y_{r_2}) \in \mathbb{R}^n$$

such that

$$|a_i| \leq (B \cdot N(\mathfrak{a}))^{1/n} \exp(D_{\sigma_i}) \text{ and } \max\{|x_j|, |y_j|\} \leq (B \cdot N(\mathfrak{a}))^{1/n} \exp(D_{\tau_j}/2)$$

for all relevant indices i, j. Clearly $\mathcal{P}'(\mathfrak{a}, B)$ contains $\mathcal{P}(\mathfrak{a}, B)$, so in order to compute $\Phi(\mathfrak{a}) \cap \mathcal{P}(\mathfrak{a}, B)$ it would suffice to compute $\Phi(\mathfrak{a}) \cap \mathcal{P}'(\mathfrak{a}, B)$ and then eliminate points that do not satisfy the inequalities defining $\mathcal{P}(\mathfrak{a}, B)$. At the expense of slightly increasing the size of the region in which we search for lattice points, we gain the advantage of having a region that is easier to work with computationally. Indeed, $\mathcal{P}'(\mathfrak{a}, B)$ is a polytope – i.e., the convex hull of a finite set of points – and the problem of enumerating lattice points in polytopes has been well studied from a theoretical as well as computational point of view [2, 3, 6].

In order to determine the points of the lattice $\Phi(\mathfrak{a})$ that lie inside the polytope $\mathcal{P}'(\mathfrak{a}, B)$, we translate this problem into one of finding *integer* lattice points in a different polytope. Let $\{\omega_1, \ldots, \omega_n\}$ be an integral basis for the ideal $\mathfrak{a}$, and let S be the $n \times n$ matrix with column vectors $\Phi(\omega_1), \ldots, \Phi(\omega_n)$. The linear isomorphism $\mathbb{R}^n \to \mathbb{R}^n$ represented by the matrix S^{-1} transforms the lattice $\Phi(\mathfrak{a})$ into the lattice $\mathbb{Z}^n$ and the polytope $\mathcal{P}'(\mathfrak{a}, B)$ into a polytope $\mathcal{X}$. Note that for any integers $s_1, \ldots, s_n$,

$$s_1\Phi(\omega_1) + \cdots + s_n\Phi(\omega_n) \in \mathcal{P}'(\mathfrak{a}, B) \iff (s_1, \ldots, s_n) \in \mathcal{X}. \tag{3.3}$$

Hence, $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ is contained in the set of all numbers x of the form $x = s_1\omega_1 + \cdots + s_n\omega_n$ with $(s_1, \ldots, s_n) \in \mathbb{Z}^n \cap \mathcal{X}$. We can therefore determine all elements of $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ once the integer points in $\mathcal{X}$ are known: indeed, it suffices to construct all numbers x of the above form and check the condition $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$. An algorithm for computing integer lattice points in polytopes is described in [6] and implemented in the software package `LattE` [1]. There is also ongoing work to include this algorithm in the `Sage` [22] software system.

To summarize this approach to computing $\mathcal{C}(\mathfrak{a}, B)$: first, a reduced basis for the lattice $\Lambda(\mathcal{O}_K^\times)$ is computed, and using it the polytope $\mathcal{P}'(\mathfrak{a}, B)$ is constructed. Computing an integral basis $\{\omega_1, \ldots, \omega_n\}$ of $\mathfrak{a}$, an $n \times n$ matrix S is defined to have columns $\Phi(\omega_1), \ldots, \Phi(\omega_n)$. The map S^{-1} is then applied to $\mathcal{P}'(\mathfrak{a}, B)$ to obtain a new polytope $\mathcal{X}$. Listing the integer points in $\mathcal{X}$ we obtain a finite list of all integer tuples $(s_1, \ldots, s_n)$ with the property that the element $x = s_1\omega_1 + \cdots + s_n\omega_n$ satisfies $\Phi(x) \in \mathcal{P}'(\mathfrak{a}, B)$. For all such elements x we then check whether $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$, thus obtaining the set $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$. Finally, we compute the norms of all elements of $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ in order to check the inequality $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a})$, and thus determine all elements of $\mathcal{C}(\mathfrak{a}, B)$.

As noted in the previous section, for the purpose of computing $\Omega(\mathfrak{a}, B)$ it would suffice to find elements representing all the orbits in $\mathcal{C}(\mathfrak{a}, B)/\{\pm 1\}$. With minor modifications, the above procedure can be used to compute only these elements instead of all $\mathcal{C}(\mathfrak{a}, B)$.

Proposition 3.8. *Let $\mathcal{H} \subset \mathbb{R}^n$ be the half-space consisting of all points whose first coordinate is non-negative. A complete set of representatives for the orbit space $\mathcal{C}(\mathfrak{a}, B)/\{\pm 1\}$ is given by all numbers of the form*

$$x = s_1\omega_1 + \cdots + s_n\omega_n$$

with $(s_1, \ldots, s_n) \in \mathbb{Z}^n \cap \mathcal{H} \cap \mathcal{X}$ satisfying $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$ and $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a})$.

Proof. Note to begin that, by definition, all elements x of the above form indeed belong to $\mathcal{C}(\mathfrak{a}, B)$. We claim that they represent all the orbits of the action of $\{\pm 1\}$. Let x be any element of $\mathcal{C}(\mathfrak{a}, B)$ and write

$$x = s_1\omega_1 + \cdots + s_n\omega_n$$

for some integers $s_1, \dots, s_n$. Replacing x with $-x$ if necessary, we may assume that $s_1 \geq 0$. Since $\Phi(x) \in \mathcal{P}'(\mathfrak{a}, B)$, then (3.3) implies that $(s_1, \dots, s_n) \in \mathcal{X}$ and hence $(s_1, \dots, s_n) \in \mathbb{Z}^n \cap \mathcal{H} \cap \mathcal{X}$. It follows that x is one of the elements listed in the statement of the proposition. □

From Proposition 3.8 and the preceding discussion we obtain the following algorithm.

Algorithm 3. Computing $\mathcal{C}(\mathfrak{a}, B)/\{\pm 1\}$

(1) Create an empty list L.
(2) Compute an LLL-reduced system of fundamental units in K.
(3) Compute the numbers D_v for all places $v|\infty$.
(4) Construct the polytope $\mathcal{P}'(\mathfrak{a}, B)$.
(5) Compute an integral basis $\omega_1, \dots, \omega_n$ for $\mathfrak{a}$.
(6) Let S be the $n \times n$ matrix with column vectors $\Phi(\omega_i)$.
(7) Construct the polytope $\mathcal{X} = S^{-1}(\mathcal{P}'(\mathfrak{a}, B))$.
(8) Find all integer lattice points in the polytope $\mathcal{H} \cap \mathcal{X}$.
(9) For all such points $(s_1, \dots, s_n)$:
  (a) Let $x = s_1\omega_1 + \cdots + s_n\omega_n$.
  (b) If $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$ and $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a})$, then include x in L.
(10) Return the list L.

One possible issue with this method for computing $\mathcal{C}(\mathfrak{a}, B)$ is that in practice the set $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ can be significantly larger than $\mathcal{C}(\mathfrak{a}, B)$, so that the step of computing the norms of all elements of $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$ is rather inefficient. We will therefore propose a different approach which reduces the number of norm computations needed. There is, however, a tradeoff between the two approaches, since the second may require a substantial number of arithmetic operations with fundamental units. If K is a number field for which these units are extremely large, it may be better to use the method described above.

3.3.2. *Second approach.* Rather than first finding all elements $x \in \mathfrak{a}$ with $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$ and then checking the condition $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a})$, we first find elements x satisfying this norm bound and then check whether $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$. In general, there will be infinitely many elements $x \in \mathfrak{a}$ with $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a})$, since any such x can be multiplied by a unit in $\mathcal{O}_K$ to obtain another element of $\mathfrak{a}$ with equal norm. However, there are only finitely many possibilities for the ideal generated by x, since there are only finitely many ideals of bounded norm in $\mathcal{O}_K$. We will show here that the set $\mathcal{C}(\mathfrak{a}, B)$ can be determined by computing a finite list of principal ideals and a finite set of units. This approach is based on ideas first introduced in [9].

Let $\mathcal{I}(\mathfrak{a}, B)$ be a set of generators for all the nonzero principal ideals that are contained in $\mathfrak{a}$ and whose norms are at most $B \cdot N(\mathfrak{a})$. We assume that distinct elements of $\mathcal{I}(\mathfrak{a}, B)$ generate distinct ideals. The elements of $\mathcal{I}(\mathfrak{a}, B)$ can be determined by using known methods for solving norm equations in number fields: applying the algorithm described in [10] (see also [11] and [19, §5.3, §6.4]) one can find generators for all principal ideals of $\mathcal{O}_K$ whose norms are of the form $k \cdot N(\mathfrak{a})$ with $1 \leq k \leq B$. Keeping only those generators that belong to $\mathfrak{a}$ we obtain $\mathcal{I}(\mathfrak{a}, B)$.

For every place $v|\infty$ we define real numbers A_v and L_v by the formulas

$$A_v = \min_{y \in \mathcal{I}(\mathfrak{a}, B)} \Lambda(y)_v,$$
$$L_v = (n_v/n)(\log B + \log N(\mathfrak{a})) + D_v - A_v.$$

Lemma 3.9. *For every nonzero element $x \in \mathcal{C}(\mathfrak{a}, B)$ there exist a unit $u \in \mathcal{O}_K^\times$ and an element $y \in \mathcal{I}(\mathfrak{a}, B)$ such that $x = uy$ and*

$$\|u\|_v \leq (B \cdot N(\mathfrak{a}))^{n_v/n} \exp(D_v - A_v) \quad \forall\, v|\infty. \tag{3.4}$$

Proof. Let $x \in \mathcal{C}(\mathfrak{a}, B)$ be any nonzero element. Since $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a})$, the ideal generated by x has norm at most $B \cdot N(\mathfrak{a})$, so x must be associate to an element of $\mathcal{I}(\mathfrak{a}, B)$. Hence, we can write $x = uy$ for some unit u and some $y \in \mathcal{I}(\mathfrak{a}, B)$. By definition, the fact that $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$ means that $\|x\|_v \leq (B \cdot N(\mathfrak{a}))^{n_v/n} \exp(D_v)$ for every place $v|\infty$. Therefore,

$$\|u\|_v = \|x\|_v \|y\|_v^{-1} \leq (B \cdot N(\mathfrak{a}))^{n_v/n} \exp(D_v) \exp(-A_v). \qquad \square$$

The above lemma shows that we can determine $\mathcal{C}(\mathfrak{a}, B)$ if we find all units u satisfying the bounds (3.4); to do this, we reduce the problem to one of finding integer points inside a polytope. Let $\mathcal{U}(\mathfrak{a}, B) \subset \mathbb{R}^r$ be the polytope consisting of all points $(t_1, \dots, t_r)$ that satisfy the inequalities $t_j \leq L_{v_j}$ for $1 \leq j \leq r$, and $t_1 + \cdots + t_r \geq -L_{v_{r+1}}$. Let $\pi : \mathbb{R}^{r+1} \to \mathbb{R}^r$ be the linear map that deletes the last coordinate, and let $\tilde{\Lambda} = \pi \circ \Lambda : K^* \to \mathbb{R}^r$. It is a standard fact that the image of the unit group $\mathcal{O}_K^\times$ under $\tilde{\Lambda}$ is a lattice of full rank in $\mathbb{R}^r$ generated by the vectors $\tilde{\Lambda}(\varepsilon_1), \dots, \tilde{\Lambda}(\varepsilon_r)$. Let T be the $r \times r$ matrix having these vectors as columns.

Lemma 3.10. *For every unit $u \in \mathcal{O}_K^\times$ satisfying the bounds* (3.4) *there exist a root of unity $\zeta \in \mu_K$ and an integer tuple $(n_1, \dots, n_r) \in T^{-1}(\mathcal{U}(\mathfrak{a}, B))$ such that $u = \zeta \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r}$.*

Proof. Let $t_1, \dots, t_r$ be the coordinates of the vector $\tilde{\Lambda}(u)$. By definition we have $t_j = \log \|u\|_{v_j}$, so (3.4) implies that $t_j \leq L_{v_j}$. Moreover, the sum of the coordinates of $\Lambda(u)$ is 0, so

$$t_1 + \cdots + t_r = -\log \|u\|_{v_{r+1}} \geq -L_{v_{r+1}},$$

and hence $\tilde{\Lambda}(u) \in \mathcal{U}(\mathfrak{a}, B)$. Write $u = \zeta \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r}$ with $\zeta \in \mu_K$ and $n_1, \dots, n_r \in \mathbb{Z}$. Then

$$n_1 \tilde{\Lambda}(\varepsilon_1) + \cdots + n_r \tilde{\Lambda}(\varepsilon_r) = \tilde{\Lambda}(u) \in \mathcal{U}(\mathfrak{a}, B).$$

Applying the linear map T^{-1} we conclude that $(n_1, \dots, n_r) \in T^{-1}(\mathcal{U}(\mathfrak{a}, B))$, and this proves the lemma. $\square$

In view of Lemma 3.10, the problem of computing $\mathcal{C}(\mathfrak{a}, B)$ is now reduced to that of finding all integer points inside the polytope $T^{-1}(\mathcal{U}(\mathfrak{a}, B))$. As mentioned earlier, this kind of problem can be solved using the algorithm developed in [6].

Putting together our results in this section we obtain the following method for computing $\mathcal{C}(\mathfrak{a}, B)$: first, the set $\mathcal{I}(\mathfrak{a}, B)$ is determined, and using this the numbers A_v and L_v are computed for every place $v|\infty$. (This requires previous knowledge of a system of fundamental units in K from which the numbers D_v are computed.) The polytope $T^{-1}(\mathcal{U}(\mathfrak{a}, B))$ is then constructed, and all integer lattice points inside it are found. For every such integer point $(n_1, \dots, n_r)$, and for every root of unity $\zeta \in \mu_K$, we then construct all numbers of the form $x = \zeta \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$ with $y \in \mathcal{I}(\mathfrak{a}, B)$. If $\Phi(x) \in \mathcal{P}(\mathfrak{a}, B)$, then we keep x because it is an element of $\mathcal{C}(\mathfrak{a}, B)$; otherwise x is discarded.

Using the above procedure we can compute the set $\mathcal{C}(\mathfrak{a}, B)$, and this could then be used to compute $\Omega(\mathfrak{a}, B)$. However, to compute $\Omega(\mathfrak{a}, B)$ using Algorithm 2 it is enough to determine representatives for the orbit space $\mathcal{C}(\mathfrak{a}, B)/\mu_K$ instead of computing all of $\mathcal{C}(\mathfrak{a}, B)$. A small change to the method described above will allow us to compute only these representatives.

Proposition 3.11. *Let $g_1, \dots, g_t$ be all the numbers of the form $g = \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$ with $y \in \mathcal{I}(\mathfrak{a}, B)$ and $(n_1, \dots, n_r) \in \mathbb{Z}^r \cap T^{-1}(\mathcal{U}(\mathfrak{a}, B))$ satisfying $\Phi(g) \in \mathcal{P}(\mathfrak{a}, B)$. Then the numbers $0, g_1, \dots, g_t$ form a complete set of representatives for the orbit space $\mathcal{C}(\mathfrak{a}, B)/\mu_K$.*

Proof. Note first of all that, by construction, the numbers $0, g_1, \dots, g_t$ all belong to $\mathcal{C}(\mathfrak{a}, B)$. Moreover, since no two elements of $\mathcal{I}(\mathfrak{a}, B)$ are associate, the orbits of $0, g_1, \dots, g_t$ are all distinct. Let $x \in \mathcal{C}(\mathfrak{a}, B)$ be nonzero. By Lemmas 3.9 and 3.10 there exist $y \in \mathcal{I}(\mathfrak{a}, B)$ and $(n_1, \dots, n_r) \in \mathbb{Z}^r \cap T^{-1}(\mathcal{U}(\mathfrak{a}, B))$ such that $x = \zeta \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$ for some $\zeta \in \mu_K$. Thus, x is in the μ_K-orbit of the element $g = \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$. Since μ_K acts

on $\mathcal{C}(\mathfrak{a}, B)$, this implies that $g \in \mathcal{C}(\mathfrak{a}, B)$, so in particular $\Phi(g) \in \mathcal{P}(\mathfrak{a}, B)$. It follows that $g \in \{g_1, \dots, g_t\}$, proving that x is in the μ_K-orbit of one of the numbers $g_1, \dots, g_t$. □

We summarize the results of this section in the following algorithm.

Algorithm 4. Computing $\mathcal{C}(\mathfrak{a}, B)/\mu_K$

(1) Create a list L containing only the element 0.
(2) Compute $\mathcal{I}(\mathfrak{a}, B)$ by solving norm equations as described above.
(3) Compute an LLL-reduced system $\{\varepsilon_1, \dots, \varepsilon_r\}$ of fundamental units in K.
(4) Compute the numbers D_v, A_v, and L_v for every place $v|\infty$.
(5) Let T be the $r \times r$ matrix with column vectors $\tilde{\Lambda}(\varepsilon_1), \dots, \tilde{\Lambda}(\varepsilon_r)$.
(6) Construct the polytope $T^{-1}(\mathcal{U}(\mathfrak{a}, B))$.
(7) Find all integer lattice points inside $T^{-1}(\mathcal{U}(\mathfrak{a}, B))$.
(8) For all such points $(n_1, \dots, n_r)$, and for every element $y \in \mathcal{I}(\mathfrak{a}, B)$:
    (a) Let $g = \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$.
    (b) If $\Phi(g) \in \mathcal{P}(\mathfrak{a}, B)$, then include g in L.
(9) Return the list L.

Algorithms 2, 3, and 4 provide two different ways of finding points of bounded height with specified ideal class. We turn now to the more general problem of determining all points of bounded height in $\mathbb{P}^N(K)$.

4. A search space for all points of bounded height

Given a real number $B \geq 1$, we wish to determine all points in the set

$$\Omega(B) = \{P \in \mathbb{P}^N(K) : H_K(P) \leq B\}.$$

If $\mathfrak{a}_1, \dots, \mathfrak{a}_h$ are ideals representing the distinct ideal classes of $\mathcal{O}_K$, then

$$\Omega(B) = \bigcup_{i=1}^{h} \Omega(\mathfrak{a}_i, B),$$

the union being disjoint. In order to compute $\Omega(B)$ it therefore suffices to compute $\Omega(\mathfrak{a}_i, B)$ for every index i, and this can be done by applying Algorithm 2. Though this approach would certainly work, there are simple modifications that can be made to shorten the computation. The crucial step for finding all the points in $\Omega(\mathfrak{a}_i, B)$ using Algorithm 2 is to compute the set $\mathcal{C}(\mathfrak{a}_i, B) \subset \mathcal{O}_K$. Thus, in the process of determining all points in $\Omega(B)$ as described above, one would compute $\mathcal{C}(\mathfrak{a}_i, B)$ for every i. In practice there can be a significant amount of overlap between the various sets $\mathcal{C}(\mathfrak{a}_1, B), \dots, \mathcal{C}(\mathfrak{a}_h, B)$, so it can happen that the same elements of $\mathcal{O}_K$ are being computed several times. In order to avoid this redundancy, we will carry out one computation of a set $\mathcal{C}(B)$ that contains all of the sets $\mathcal{C}(\mathfrak{a}_i, B)$, and then for each i the elements of $\mathcal{C}(\mathfrak{a}_i, B)$ will be found by searching through $\mathcal{C}(B)$.

Let $\mathbf{N} = \max_i N(\mathfrak{a}_i)$ and let $\mathcal{P}(B)$ be the subset of $\prod_{v|\infty} K_v \cong \mathbb{R}^n$ consisting of all points (s_v) such that

$$|s_v| \leq (B \cdot \mathbf{N})^{1/n} \exp(D_v/n_v) \quad \forall\, v|\infty.$$

Note that $\mathcal{P}(B)$ contains $\mathcal{P}(\mathfrak{a}_i, B)$ for every i. Define a set $\mathcal{N}$ of non-negative integers by

$$\mathcal{N} = \{0\} \cup \bigcup_{i=1}^{h} \{k \cdot N(\mathfrak{a}_i) : 1 \leq k \leq B\},$$

and let

$$\mathcal{C}(B) = \{x \in \mathcal{O}_K : \Phi(x) \in \mathcal{P}(B) \text{ and } |N_{K/\mathbb{Q}}(x)| \in \mathcal{N}\},$$

so that $\mathcal{C}(B)$ contains $\mathcal{C}(\mathfrak{a}_i, B)$ for every i. The methods of §3.3 can be easily adapted to compute $\mathcal{C}(B)$. Once this has been done, the various sets $\mathcal{C}(\mathfrak{a}_i, B)$ can be determined by checking, for every element $x \in \mathcal{C}(B)$, whether x satisfies the conditions $\Phi(x) \in \mathcal{P}(\mathfrak{a}_i, B)$ and $|N_{K/\mathbb{Q}}(x)| \leq B \cdot N(\mathfrak{a}_i)$. The sets $\mathcal{C}(\mathfrak{a}_1, B), \ldots, \mathcal{C}(\mathfrak{a}_h, B)$ are thus obtained, and can then be used to determine all points in the sets $\Omega(\mathfrak{a}_i, B)$.

Now, in order to compute $\Omega(\mathfrak{a}_i, B)$ using Algorithm 2 it suffices to find representatives for the orbit spaces $\mathcal{C}(\mathfrak{a}_i, B)/\mu_K$ or $\mathcal{C}(\mathfrak{a}_i, B)/\{\pm 1\}$, so it would be desirable that group actions on $\mathcal{C}(B)$ could be used to find these representatives instead of computing the entire set $\mathcal{C}(\mathfrak{a}_i, B)$. This can indeed be done without any additional work: the groups μ_K and $\{\pm 1\}$ act on $\mathcal{C}(B)$, and with minor changes the methods of §3.3 can be used to compute representatives for the orbits of these actions; the details of this are discussed below. By selecting the elements that belong to $\mathcal{C}(\mathfrak{a}_i, B)$ for each i we obtain representatives for the orbit spaces $\mathcal{C}(\mathfrak{a}_i, B)/\mu_K$ or $\mathcal{C}(\mathfrak{a}_i, B)/\{\pm 1\}$. These representatives can then be used in Algorithm 2 to compute $\Omega(\mathfrak{a}_i, B)$.

Depending on which orbit space is computed, $\mathcal{C}(B)/\{\pm 1\}$ or $\mathcal{C}(B)/\mu_K$, our discussion above yields a different method to compute $\Omega(B)$. We will henceforth denote by M1 the algorithm that uses the action of $\{\pm 1\}$, and by M2 the algorithm that uses the action of μ_K.

For the algorithm M1, the methods of §3.3.1 should be applied to compute $\mathcal{C}(B)/\{\pm 1\}$. Let $\mathcal{P}'(B) \subset \mathbb{R}^n$ be the region consisting of all points

$$(a_1, \ldots, a_{r_1}; x_1, y_1, \ldots, x_{r_2}, y_{r_2}) \in \mathbb{R}^n$$

such that

$$|a_i| \leq (B \cdot \mathbf{N})^{1/n} \exp(D_{\sigma_i}) \quad \text{and} \quad \max\{|x_j|, |y_j|\} \leq (B \cdot \mathbf{N})^{1/n} \exp(D_{\tau_j}/2)$$

for all indices $1 \leq i \leq r_1$ and $1 \leq j \leq r_2$. Note that $\mathcal{P}'(B)$ contains $\mathcal{P}(B)$. Let $\{\omega_1, \ldots, \omega_n\}$ be an integral basis for $\mathcal{O}_K$, and let S be the $n \times n$ matrix with column vectors $\Phi(\omega_1), \ldots, \Phi(\omega_n)$. Finally, let $\mathcal{X}$ be the polytope $S^{-1}(\mathcal{P}'(B))$. A minor modification of the proof of Proposition 3.8 yields the following result.

Proposition 4.1. *Let $\mathcal{H} \subset \mathbb{R}^n$ be the half-space consisting of all points whose first coordinate is non-negative. A complete set of representatives for the orbit space $\mathcal{C}(B)/\{\pm 1\}$ is given by all numbers of the form*

$$x = s_1\omega_1 + \cdots + s_n\omega_n$$

with $(s_1, \ldots, s_n) \in \mathbb{Z}^n \cap \mathcal{H} \cap \mathcal{X}$ satisfying $\Phi(x) \in \mathcal{P}(B)$ and $|N_{K/\mathbb{Q}}(x)| \in \mathcal{N}$.

The above proposition suggests an algorithm for computing $\mathcal{C}(B)/\{\pm 1\}$ that is analogous (in fact, nearly identical) to Algorithm 3.

From our work up to this point we obtain the following description of M1. Briefly, what the algorithm below does is to compute $\mathcal{C}(B)/\{\pm 1\}$, then intersect with each set $\mathcal{C}(\mathfrak{a}_i, B)$ to determine $\mathcal{C}(\mathfrak{a}_i, B)/\{\pm 1\}$, and finally use Algorithm 2 to compute each set $\Omega(\mathfrak{a}_i, B)$.

Algorithm 5 (M1)**.** Computing $\Omega(B)$ using the action of $\{\pm 1\}$.

(1) Create an empty list $\mathcal{L}$. This list will store the points belonging to $\Omega(B)$.
(2) Compute an integral basis $\omega_1, \ldots, \omega_n$ for $\mathcal{O}_K$.
(3) Determine ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_h$ representing the distinct ideal classes of $\mathcal{O}_K$.
(4) Compute an LLL-reduced system of fundamental units in $\mathcal{O}_K$.
(5) Construct the set $\mathcal{N}$.
(6) Compute the numbers D_v for all places $v|\infty$.
(7) Construct the polytope $\mathcal{P}'(B)$.
(8) Let S be the $n \times n$ matrix with column vectors $\Phi(\omega_i)$.
(9) Construct the polytope $\mathcal{X} = S^{-1}(\mathcal{P}'(B))$.
(10) Create an empty set L. This set will store representatives for $\mathcal{C}(B)/\{\pm 1\}$.
(11) Find all integer lattice points in the polytope $\mathcal{H} \cap \mathcal{X}$.

(12) For all such points $(s_1, \ldots, s_n)$:
  (a) Let $x = s_1\omega_1 + \cdots + s_n\omega_n$.
  (b) If $\Phi(x) \in \mathcal{P}(B)$ and $|N_{K/\mathbb{Q}}(x)| \in \mathcal{N}$, then include x in L.
(13) For each ideal $\mathfrak{a} \in \{\mathfrak{a}_1, \ldots, \mathfrak{a}_h\}$:
  (a) Fix an ordering $g_0, g_1, \ldots, g_t$ of the elements of $L \cap \mathcal{C}(\mathfrak{a}, B)$.
  (b) For every tuple of indices $(i_0, \ldots, i_N)$ such that $0 \le i_0 \le i_1 \le \cdots \le i_N \le t$:
    (i) Let X be the point $(g_{i_0}, \ldots, g_{i_N})$.
    (ii) If $\mathfrak{a} = (g_{i_0}, \ldots, g_{i_N})$ and $H_\infty(X) \le B \cdot N(\mathfrak{a})$, then:
    Include in $\mathcal{L}$ all the points in $\mathcal{O}(X)$.
(14) Return the list $\mathcal{L}$.

For the algorithm M2, the methods of §3.3.2 should be applied to compute $\mathcal{C}(B)/\mu_K$. Let $\mathcal{I}(B)$ be a set of generators for all the nonzero principal ideals of $\mathcal{O}_K$ whose norms are in $\mathcal{N}$. For every place $v|\infty$, define real numbers

$$a_v = \min_{y \in \mathcal{I}(B)} \Lambda(y)_v,$$
$$\ell_v = (n_v/n)(\log B + \log \mathbf{N}) + D_v - a_v.$$

Let $\mathcal{U}(B) \subset \mathbb{R}^r$ be the polytope consisting of all points $(t_1, \ldots, t_r)$ that satisfy the inequalities $t_j \le \ell_{v_j}$ for $1 \le j \le r$, and $t_1 + \cdots + t_r \ge -\ell_{v_{r+1}}$. Finally, let T be the $r \times r$ matrix with column vectors $\tilde{\Lambda}(\varepsilon_1), \ldots, \tilde{\Lambda}(\varepsilon_r)$. The results of §3.3.2 can be modified in an obvious way to obtain the following result.

Proposition 4.2. *Let $g_1, \ldots, g_t$ be all the numbers of the form $g = \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$ with $y \in \mathcal{I}(B)$ and $(n_1, \ldots, n_r) \in \mathbb{Z}^r \cap T^{-1}(\mathcal{U}(B))$ satisfying $\Phi(g) \in \mathcal{P}(B)$. Then the numbers $0, g_1, \ldots, g_t$ form a complete set of representatives for the orbit space $\mathcal{C}(B)/\mu_K$.*

From the above proposition we deduce an algorithm for computing $\mathcal{C}(B)/\mu_K$ that is analogous to Algorithm 4. We include this algorithm in the following complete description of M2.

Algorithm 6 (M2)**.** Computing $\Omega(B)$ using the action of μ_K.

(1) Create an empty list $\mathcal{L}$. This list will store the points belonging to $\Omega(B)$.
(2) Determine ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_h$ representing the distinct ideal classes of $\mathcal{O}_K$.
(3) Compute an LLL-reduced system $\{\varepsilon_1, \ldots, \varepsilon_r\}$ of fundamental units in K.
(4) Construct the set $\mathcal{N}$.
(5) Compute the set $\mathcal{I}(B)$ by solving norm equations.
(6) Compute the numbers D_v, a_v, and ℓ_v for every place $v|\infty$.
(7) Let T be the $r \times r$ matrix with column vectors $\tilde{\Lambda}(\varepsilon_1), \ldots, \tilde{\Lambda}(\varepsilon_r)$.
(8) Construct the polytope $T^{-1}(\mathcal{U}(B))$.
(9) Create the set $L = \{0\}$. This set will store representatives for $\mathcal{C}(B)/\mu_K$.
(10) Find all integer lattice points inside $T^{-1}(\mathcal{U}(B))$.
(11) For all such points $(n_1, \ldots, n_r)$, and for every element $y \in \mathcal{I}(B)$:
  (a) Let $g = \varepsilon_1^{n_1} \cdots \varepsilon_r^{n_r} y$.
  (b) If $\Phi(g) \in \mathcal{P}(B)$, then include g in L.
(12) For each ideal $\mathfrak{a} \in \{\mathfrak{a}_1, \ldots, \mathfrak{a}_h\}$:
  (a) Fix an ordering $g_0, g_1, \ldots, g_t$ of the elements of $L \cap \mathcal{C}(\mathfrak{a}, B)$.
  (b) For every tuple of indices $(i_0, \ldots, i_N)$ such that $0 \le i_0 \le i_1 \le \cdots \le i_N \le t$:
    (i) Let X be the point $(g_{i_0}, \ldots, g_{i_N})$.
    (ii) If $\mathfrak{a} = (g_{i_0}, \ldots, g_{i_N})$ and $H_\infty(X) \le B \cdot N(\mathfrak{a})$, then:
    Include in $\mathcal{L}$ all the points in $\mathcal{O}(X)$.

(13) Return the list $\mathcal{L}$.

To conclude our discussion we mention one optional modification that in many cases leads to an improvement in the performance of both M1 and M2. For any point $P = [x_0, \dots, x_N] \in \mathbb{P}^N(K)$ and any automorphism $\sigma \in \mathrm{Aut}(K/\mathbb{Q})$, we denote by P^σ the point

$$P^\sigma = [\sigma(x_0), \dots, \sigma(x_N)] \in \mathbb{P}^N(K).$$

Note that P and P^σ have the same height and have Galois-conjugate ideal classes. Using this observation one sees that if $\mathfrak{c}_1, \dots, \mathfrak{c}_s \in \{\mathfrak{a}_1, \dots, \mathfrak{a}_h\}$ are ideals representing the distinct orbits of the action of the group $G = \mathrm{Aut}(K/\mathbb{Q})$ on the ideal class group $\mathrm{Cl}(\mathcal{O}_K)$, then

$$\Omega(B) = \bigcup_{i=1}^{s} G \cdot \Omega(\mathfrak{c}_i, B).$$

Thus, in order to compute $\Omega(B)$ it suffices to compute the sets $\Omega(\mathfrak{c}_i, B)$ for $i = 1, \dots, s$ and then let the group G act on these sets. If s is substantially smaller than h, then the reduction in the number of ideals $\mathfrak{a}$ for which the set $\Omega(\mathfrak{a}, B)$ is computed can easily make up for the additional cost of dividing the class group into G-orbits.

5. The case of unit rank zero

We discuss here the problem of enumerating all points in the set

$$\Omega(B) = \{P \in \mathbb{P}^N(K) : H_K(P) \le B\}$$

when K is a number field with finite unit group (i.e., $K = \mathbb{Q}$ or an imaginary quadratic field). The reasons for treating this case separately are twofold: first, the height function takes a particularly simple shape in this case, reducing the problem to a computation of elements of $\mathcal{O}_K$ with bounded norm; second, for the computation of ideal class groups of imaginary quadratic fields there are specialized algorithms that are faster than the methods that apply to arbitrary number fields [4, 12, 13].

A method for listing all points in $\Omega(B)$ when $K = \mathbb{Q}$ is easily deduced from the following proposition. We will be using here the notation introduced in §3.2.

Proposition 5.1. *Let $\mathcal{M}$ denote the set of all tuples $(x_0, \dots, x_N) \in \mathbb{Z}^{N+1}$ such that $0 \le x_0 \le \cdots \le x_N \le B$ and $\gcd(x_0, \dots, x_N) = 1$. Then*

$$\{P \in \mathbb{P}^N(\mathbb{Q}) : H_{\mathbb{Q}}(P) \le B\} = \bigcup_{X \in \mathcal{M}} \mathcal{O}(X).$$

Proof. This is a simple exercise using the fact that if $P \in \mathbb{P}^N(\mathbb{Q})$ satisfies $H_{\mathbb{Q}}(P) \le B$, then it is possible to write $P = [a_0, \dots, a_N]$ with $a_i \in \mathbb{Z} \cap [-B, B]$ for all i and $\gcd(a_0, \dots, a_N) = 1$. □

We assume henceforth that K is an imaginary quadratic field. As has been noted earlier, in order to find all points of bounded height in $\mathbb{P}^N(K)$ it suffices to find all points of bounded height with given ideal class. For any nonzero ideal $\mathfrak{a}$ of $\mathcal{O}_K$ and real number $B \ge 1$, define

$$\Omega(\mathfrak{a}, B) = \{P \in \mathbb{P}^N(K) : \mathrm{Cl}(P) = \mathrm{Cl}(\mathfrak{a}) \text{ and } H_K(P) \le B\}.$$

The following proposition reduces the problem of computing $\Omega(\mathfrak{a}, B)$ to the computation of the set

$$\mathcal{C}(\mathfrak{a}, B) = \{\gamma \in \mathfrak{a} : N_{K/\mathbb{Q}}(\gamma) \le B \cdot N(\mathfrak{a})\}.$$

Proposition 5.2. *Let $g_0, g_1, \ldots, g_t \in \mathcal{C}(\mathfrak{a}, B)$ be elements representing all the orbits of the action of μ_K on $\mathcal{C}(\mathfrak{a}, B)$. Let $\mathcal{M}$ be the set of all tuples $X = (g_{i_0}, \ldots, g_{i_N})$ such that $0 \leq i_0 \leq i_1 \leq \cdots \leq i_N \leq t$ and $\mathfrak{a} = (g_{i_0}, \ldots, g_{i_N})$. Then*

$$\Omega(\mathfrak{a}, B) = \bigcup_{X \in \mathcal{M}} \mathcal{O}(X).$$

Proof. Let $\sigma, \overline{\sigma}$ be the embeddings $K \hookrightarrow \mathbb{C}$. For any element $\alpha \in K$ we have $N_{K/\mathbb{Q}}(\alpha) = \sigma(\alpha)\overline{\sigma}(\alpha) = |\sigma(\alpha)|^2$. Hence, for any tuple $(x_0, \ldots, x_N) \in K^{N+1}$ we obtain

$$H_\infty(x_0, \ldots, x_N) = \max\{N_{K/\mathbb{Q}}(x_0), \ldots, N_{K/\mathbb{Q}}(x_N)\}. \tag{5.1}$$

Suppose that $P \in \Omega(\mathfrak{a}, B)$. Since $\mathrm{Cl}(P) = \mathrm{Cl}(\mathfrak{a})$, there are homogeneous coordinates $[x_0, \ldots, x_N]$ for P such that $\mathfrak{a} = (x_0, \ldots, x_N)$. The condition $H_K(P) \leq B$ is then equivalent to $H_\infty(x_0, \ldots, x_N) \leq B \cdot N(\mathfrak{a})$, so by (5.1) we see that $x_i \in \mathcal{C}(\mathfrak{a}, B)$ for every i. The remainder of the proof is entirely analogous to the proof of Proposition 3.5. □

Let $\mathfrak{a}_1, \ldots, \mathfrak{a}_h$ be ideals representing the ideal classes of $\mathcal{O}_K$. If we can determine representatives for each orbit space $\mathcal{C}(\mathfrak{a}_i, B)/\mu_K$, then Proposition 5.2 can be used compute $\Omega(\mathfrak{a}_i, B)$ for each i, and thus $\Omega(B)$ is obtained. Though it is possible to compute each set $\mathcal{C}(\mathfrak{a}_i, B)$ separately, it would be more efficient to apply some of the ideas introduced in §4; in particular, we should compute one set $\mathcal{C}(B)$ that contains all of the sets $\mathcal{C}(\mathfrak{a}_i, B)$. Define

$$\mathcal{N} = \{0\} \cup \bigcup_{i=1}^{h} \{k \cdot N(\mathfrak{a}_i) : 1 \leq k \leq B\},$$

and let

$$\mathcal{C}(B) = \{\gamma \in \mathcal{O}_K : N_{K/\mathbb{Q}}(\gamma) \in \mathcal{N}\},$$

so that $\mathcal{C}(B)$ contains every set $\mathcal{C}(\mathfrak{a}_i, B)$. Note that the group μ_K acts on $\mathcal{C}(B)$. A set L of representatives of the orbits of this action can be found by solving norm equations in K, for instance using the algorithm given in [10]. Having done this, representatives of $\mathcal{C}(\mathfrak{a}_i, B)/\mu_K$ can be obtained for every index i by intersecting L with $\mathcal{C}(\mathfrak{a}_i, B)$. We are thus led to the following algorithm.

Algorithm 7. Computing $\Omega(B)$ when K is imaginary quadratic.

(1) Create an empty list $\mathcal{L}$. This list will store the points belonging to $\Omega(B)$.
(2) Determine ideals $\mathfrak{a}_1, \ldots, \mathfrak{a}_h$ representing the ideal classes of $\mathcal{O}_K$.
(3) Construct the set $\mathcal{N}$.
(4) Create an empty set L. This set will store representatives for $\mathcal{C}(B)/\mu_K$.
(5) For every number $m \in \mathcal{N}$:
    Include in L all elements of $\mathcal{O}_K$ with norm m, modulo units.
(6) For each ideal $\mathfrak{a} \in \{\mathfrak{a}_1, \ldots, \mathfrak{a}_h\}$:
    (a) Fix an ordering $g_0, g_1, \ldots, g_t$ of the elements of the set $L \cap \mathcal{C}(\mathfrak{a}, B)$.
    (b) For every tuple of indices $(i_0, \ldots, i_N)$ such that $0 \leq i_0 \leq i_1 \leq \cdots \leq i_N \leq t$:
        If $\mathfrak{a} = (g_{i_0}, \ldots, g_{i_N})$, then include in $\mathcal{L}$ all the points in $\mathcal{O}(X)$.
(7) Return the list $\mathcal{L}$.

As in the case of number fields with positive unit rank, Algorithm 7 can be improved by using the action of the group $\mathrm{Gal}(K/\mathbb{Q})$ on the ideal class group of $\mathcal{O}_K$. See the final paragraphs of §4 for details.

6. Efficiency of the algorithm

In this section we carry out an analysis of the efficiency of our method to compute points of bounded height in $\mathbb{P}^N(K)$. Since the problem of computing $\Omega(B)$ is reduced to computing sets $\Omega(\mathfrak{a}, B)$ for a finite list of ideals $\mathfrak{a}$, we focus first on determining how efficient our method for computing such sets is. For this analysis we use Algorithm 1, which is the simplest description of our method.

As a measure of the efficiency of the algorithm we consider how many points it generates and how this quantity compares to the size of $\Omega(\mathfrak{a}, B)$. We define the *search space* of Algorithm 1 to be the set

$$\mathcal{S}(\mathfrak{a}, B) = \{[x_0, \dots, x_N] \in \mathbb{P}^N(K) : x_i \in \mathcal{C}(\mathfrak{a}, B) \text{ for all } i\}. \tag{6.1}$$

The elements of $\mathcal{S}(\mathfrak{a}, B)$ are all the points that would be considered in the algorithm while searching for points in $\Omega(\mathfrak{a}, B)$. Since the size of the search space is largely determined by the size of $\mathcal{C}(\mathfrak{a}, B)$, we begin by understanding the latter. We will bound the size of $\mathcal{C}(\mathfrak{a}, B)$ by the size of the set

$$\widetilde{\mathcal{C}}(\mathfrak{a}, B) = \{x \in \mathfrak{a} : \Phi(x) \in \mathcal{P}(\mathfrak{a}, B)\} \supseteq \mathcal{C}(\mathfrak{a}, B).$$

Proposition 6.1. *There is a constant c, depending only on K and the choice of fundamental units, such that*

$$\#\widetilde{\mathcal{C}}(\mathfrak{a}, B) = cB + O(B^{1-1/n}).$$

Proof. We need the following classical result concerning lattice points in homogeneously expanding domains; see [14, Chap. VI, §2] for details.

Lemma 6.2. *Let D be a subset of $\mathbb{R}^n$ and L a lattice in $\mathbb{R}^n$ with fundamental domain F. Assume that the boundary of D is $(n-1)$-Lipschitz parametrizable. Then the number of lattice points in tD, for $t \in \mathbb{R}_{>0}$, satisfies*

$$\#(L \cap tD) = \frac{\operatorname{Vol} D}{\operatorname{Vol} F} \cdot t^n + O(t^{n-1}).$$

To prove the proposition we apply the above lemma to the lattice $L = \Phi(\mathfrak{a})$ and the domain $D = \mathcal{P}(\mathfrak{a}, 1)$. One can see from the definitions that $\mathcal{P}(\mathfrak{a}, B) = B^{1/n}\mathcal{P}(\mathfrak{a}, 1)$; hence, applying the lemma we obtain

$$\begin{aligned} \#\widetilde{\mathcal{C}}(\mathfrak{a}, B) &= \#\left(\Phi(\mathfrak{a}) \cap \mathcal{P}(\mathfrak{a}, B)\right) = \#\left(\Phi(\mathfrak{a}) \cap B^{1/n}\mathcal{P}(\mathfrak{a}, 1)\right) \\ &= \frac{\operatorname{Vol} \mathcal{P}(\mathfrak{a}, 1)}{\operatorname{Vol} F(\mathfrak{a})} \cdot B + O(B^{1-1/n}), \end{aligned} \tag{6.2}$$

where $F(\mathfrak{a})$ is the fundamental parallelotope defined in (2.4). The region $\mathcal{P}(\mathfrak{a}, 1) \subset \mathbb{R}^n$ is a cartesian product of intervals of length $2 \cdot N(\mathfrak{a})^{1/n} \exp(D_v)$, where v ranges over the real places of K, and disks of radius $N(\mathfrak{a})^{1/n} \exp(D_v/2)$, where v ranges over the complex places. Thus,

$$\begin{aligned} \operatorname{Vol} \mathcal{P}(\mathfrak{a}, 1) &= \left(\prod_{v \text{ real}} 2 \cdot N(\mathfrak{a})^{1/n} \exp(D_v)\right) \cdot \left(\prod_{v \text{ complex}} \pi \cdot N(\mathfrak{a})^{2/n} \exp(D_v)\right) \\ &= 2^{r_1} \pi^{r_2} N(\mathfrak{a}) \exp(\textstyle\sum_v D_v). \end{aligned}$$

Since $\operatorname{Vol} F(\mathfrak{a}) = 2^{-r_2} |\Delta_K|^{1/2} N(\mathfrak{a})$, the coefficient of B in (6.2) is the constant c given by

$$c = \frac{2^{r_1}(2\pi)^{r_2} \exp(\sum_v D_v)}{|\Delta_K|^{1/2}}. \tag{6.3}$$

This proves the proposition, since c depends only on standard invariants of K and on the numbers D_v, which are determined by the choice of fundamental units. □

Having proved asymptotic bounds on the size of $\widetilde{\mathcal{C}}(\mathfrak{a}, B)$, we can now bound the size of the search space $\mathcal{S}(\mathfrak{a}, B)$.

Proposition 6.3. *Let $m = N + 1$ and define a constant C by*

$$C = \frac{w \cdot \zeta_K(m) \exp(m \sum_v D_v)}{R \cdot m^r}, \tag{6.4}$$

where w is the number of roots of unity in K; R is the regulator; r is the rank of the unit group; and ζ_K is the zeta function. Then

$$\limsup_{B \to \infty} \frac{\#\mathcal{S}(\mathfrak{a}, B)}{\#\Omega(\mathfrak{a}, B)} \leq C.$$

Proof. We have $\#\mathcal{S}(\mathfrak{a}, B) \leq (\#\mathcal{C}(\mathfrak{a}, B))^m \leq (\#\widetilde{\mathcal{C}}(\mathfrak{a}, B))^m$, so Proposition 6.1 implies that

$$\#\mathcal{S}(\mathfrak{a}, B) \leq c^m B^m + O(B^{m-1/n}). \tag{6.5}$$

Now, Schanuel [20] showed that

$$\#\Omega(\mathfrak{a}, B) = B^m \cdot m^r \cdot \frac{R/w}{\zeta_K(m)} \left(\frac{2^{r_1}(2\pi)^{r_2}}{|\Delta_K|^{1/2}} \right)^m + O(B^{m-1/n}). \tag{6.6}$$

The result follows by dividing (6.5) and (6.6) and letting $B \to \infty$. □

We can now prove our main result concerning the size of the search space of our method for computing $\Omega(B)$. If $\mathfrak{a}_1, \ldots, \mathfrak{a}_h$ are ideals representing the distinct ideal classes of $\mathcal{O}_K$, then

$$\Omega(B) = \bigcup_{i=1}^{h} \Omega(\mathfrak{a}_i, B),$$

so $\Omega(B)$ can be obtained by computing $\Omega(\mathfrak{a}_i, B)$ for every index i. With sets $\mathcal{S}(\mathfrak{a}_i, B)$ defined as in (6.1), the search space of this method is the set

$$\mathcal{S}(B) = \bigcup_{i=1}^{h} \mathcal{S}(\mathfrak{a}_i, B),$$

whose size we now compare to that of $\Omega(B)$.

Theorem 6.4. *Let h be the class number of $\mathcal{O}_K$ and let C be the constant defined in* (6.4)*. Then*

$$\limsup_{B \to \infty} \frac{\#\mathcal{S}(B)}{\#\Omega(B)} \leq hC.$$

Proof. From the definitions it follows that

$$\frac{\#\mathcal{S}(B)}{\#\Omega(B)} \leq \sum_{i=1}^{h} \frac{\#\mathcal{S}(\mathfrak{a}_i, B)}{\#\Omega(B)} \leq \sum_{i=1}^{h} \frac{\#\mathcal{S}(\mathfrak{a}_i, B)}{\#\Omega(\mathfrak{a}_i, B)}.$$

The theorem is then a consequence of Proposition 6.3. □

7. Sample Computations

The main algorithms of this paper, namely M1 and M2 (Algorithms 5 and 6), have been implemented using `Sage` [22]. We give below a series of computations of points of small height in projective spaces using these algorithms. Since our methods apply in particular to the space $\mathbb{P}^1$, and listing points of bounded height in $\mathbb{P}^1$ is equivalent to listing elements of K with bounded height, we also compare our methods to the algorithm developed in [9] for finding elements of small height in number fields. All computations have been done on a MacBook Pro with a 2.7 GHz processor and 16 GB of memory.

7.1. **Comparison of M1 and M2.** The three tables shown below summarize the results of computations of points of bounded height in $\mathbb{P}^N(K)$ for three number fields K and for $N = 1, 2, 3$. In all cases, the height bound B was taken to be 20. The purpose of these examples is to compare the running times of M1 and M2 when applied over number fields of various degrees and projective spaces of various dimensions. The results suggest that M2 is a significantly better method than M1, and indeed we have not found any examples – among many computations – where M1 performs better than M2. Theoretically, M1 has an advantage over M2 in that it does not require arithmetic operations with fundamental units; however, this on its own does not seem to make M1 a faster method. The difference in running times is largely due to the fact that M2 generates all points in $\Omega(B)$ starting from the set $\mathcal{C}(\mathfrak{a}, B)/\mu_K$ rather than the larger set $\mathcal{C}(\mathfrak{a}, B)/\{\pm 1\}$.

N	M1 time	M2 time	Points found
1	0.98 s	0.26 s	504
2	39 s	10 s	20,401
3	2,465 s	513 s	607,344

TABLE 1. Computations over the field $K = \mathbb{Q}(\sqrt{17})$

N	M1 time	M2 time	Points found
1	0.43 s	0.27 s	452
2	25 s	14 s	23,725
3	2,707	889 s	888,872

TABLE 2. Computations over the field $\mathbb{Q}(\sqrt[3]{2})$

N	M1 time	M2 time	Points found
1	11 s	0.41 s	842
2	10,483 s	51 s	72,091
3	-	10,407 s	4,926,644

TABLE 3. Computations over the field $\mathbb{Q}(\sqrt[4]{-1})$

7.2. **Elements of bounded height in number fields.** Let K be a number field and let $x \in K$. The height of x is defined to be the number $H_K(x) = H_K(P)$, where $P = [x, 1] \in \mathbb{P}^1(K)$. The problem of listing all elements of K up to a given height bound was studied in the article [9], where an algorithm – denoted here by DK – was developed to solve this problem. Since the methods M1 and M2 can be applied to list all elements of bounded height in $\mathbb{P}^1(K)$, it is natural to wonder which algorithm (M1, M2, or DK) is fastest in this context. Given that M1 appears to be generally slower than M2, we will compare only M2 and DK.

The four tables shown below list the running times for several computations done with M2 and DK over quadratic, cubic, and quartic number fields, and with several different height bounds B. In the tables, the underlying number field K is either given explicitly as an extension of $\mathbb{Q}$, or a defining polynomial for it is given in the variable x. We will denote by g a root of the defining polynomial of K, so that $K = \mathbb{Q}(g)$.

As indicated by these computations, neither algorithm is always better than the other, at least in their current implementations. There are examples where M2 is significantly faster than DK (see tables 4 and 5), and there are cases where the opposite happens (see tables 6 and 7). This phenomenon appears to be tied to the size of the fundamental units chosen for K: when the units are relatively small, DK performs better than

M2, but when the units are fairly large, M2 is faster. The fields K used for the computations in the tables below illustrate the effect of the size of the fundamental units on DK and M2. All of these fields have unit rank 1. Fundamental units for the fields in Tables 4 and 5, where M2 performs best, are given by

$$1077170348190210679608465 2 \cdot g - 11968230284425768995908496 41$$

and

$$175971701235126787623 61 \cdot g^2 + 14942822416894246257476 66 \cdot g + 70844652623253460553144 39,$$

respectively. In contrast, fundamental units for the fields in Tables 6 and 7, where DK performs best, are given by

$$28g + 295 \quad \text{and} \quad g^3 + 2g^2 + 2g - 2.$$

B	DK time	M2 time	Elements found
100	1.25 s	0.76 s	479
1,000	92 s	40 s	73,007
5,000	3,358 s	1,299 s	1,826,367

TABLE 4. Elements of bounded height in $\mathbb{Q}(\sqrt{12345})$

B	DK time	M2 time	Elements found
100	0.82 s	0.51 s	263
1,000	55 s	19 s	27,603
5,000	1,682 s	475 s	731,755

TABLE 5. Elements of bounded height in $K : x^3 - x + 123$

B	DK time	M2 time	Elements found
100	1.32 s	1.97 s	2,875
1,000	49 s	96 s	275,615
5,000	1,402 s	3,019 s	6,795,587

TABLE 6. Elements of bounded height in $\mathbb{Q}(\sqrt{111})$

B	DK time	M2 time	Elements found
100	0.69 s	0.82 s	299
1,000	13 s	23 s	42,067
5,000	202 s	431 s	1,092,203

TABLE 7. Elements of bounded height in $K : x^4 - x + 11$

References

[1] V. Baldoni et al. *A User's Guide for LattE integrale v1.6, 2013*. Available online at http://www.math.ucdavis.edu/~la

[2] Alexander Barvinok and James E. Pommersheim. "An algorithmic theory of lattice points in polyhedra". *New perspectives in algebraic combinatorics (Berkeley, CA, 1996–97)*. Vol. 38. Math. Sci. Res. Inst. Publ. Cambridge Univ. Press, 1999, pp. 91–147.

[3] Alexander I. Barvinok. "A Polynomial Time Algorithm for Counting Integral Points in Polyhedra When the Dimension Is Fixed". *Mathematics of Operations Research* 19.4 (1994), pp. 769–779.

[4] Jean-François Biasse. "Improvements in the computation of ideal class groups of imaginary quadratic number fields". *Adv. Math. Commun.* 4.2 (2010), pp. 141–154.

[5] Henri Cohen. *A course in computational algebraic number theory*. Vol. 138. Graduate Texts in Mathematics. Springer-Verlag, 1993.

[6] Jesús A. De Loera et al. "Effective lattice point counting in rational convex polytopes". *J. Symbolic Comput.* 38.4 (2004), pp. 1273–1302.

[7] Lassina Dembélé and Abhinav Kumar. *Examples of abelian surfaces with everywhere good reduction*. Preprint, http://arxiv.org/abs/1309.3821.

[8] John R. Doyle, Xander Faber, and David Krumm. *Preperiodic points for quadratic polynomials over quadratic fields*. Preprint, http://arxiv.org/abs/1309.6401.

[9] John R. Doyle and David Krumm. "Computing algebraic numbers of bounded height". *Math. Comp.* (To appear).

[10] U. Fincke and M. Pohst. "A procedure for determining algebraic integers of given norm". *Computer algebra (London, 1983)*. Vol. 162. Lecture Notes in Comput. Sci. Springer, 1983, pp. 194–202.

[11] U. Fincke and M. Pohst. "Improved methods for calculating vectors of short length in a lattice, including a complexity analysis". *Math. Comp.* 44.170 (1985), pp. 463–471.

[12] James L. Hafner and Kevin S. McCurley. "A rigorous subexponential algorithm for computation of class groups". *J. Amer. Math. Soc.* 2.4 (1989), pp. 837–850.

[13] Michael J. Jacobson Jr. "Applying sieving to the computation of quadratic class groups". *Math. Comp.* 68.226 (1999), pp. 859–867.

[14] Serge Lang. *Algebraic number theory*. Second. Vol. 110. Graduate Texts in Mathematics. Springer-Verlag, 1994.

[15] Serge Lang. *Fundamentals of Diophantine geometry*. New York: Springer-Verlag, 1983.

[16] A. K. Lenstra, H. W. Lenstra Jr., and L. Lovász. "Factoring polynomials with rational coefficients". *Math. Ann.* 261.4 (1982), pp. 515–534.

[17] Phong Q. Nguyen and Damien Stehlé. "An LLL algorithm with quadratic complexity". *SIAM J. Comput.* 39.3 (2009), pp. 874–903.

[18] A. Pethő and S. Schmitt. "Elements with bounded height in number fields". *Period. Math. Hungar.* 43.1-2 (2001), pp. 31–41.

[19] M. Pohst and H. Zassenhaus. *Algorithmic algebraic number theory*. Vol. 30. Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1997.

[20] Stephen Hoel Schanuel. "Heights in number fields". *Bull. Soc. Math. France* 107.4 (1979), pp. 433–449.

[21] Joseph H. Silverman. *The arithmetic of dynamical systems*. Vol. 241. Graduate Texts in Mathematics. Springer, 2007.

[22] W. A. Stein et al. *Sage Mathematics Software (Version 5.9)*. http://www.sagemath.org. The Sage Development Team. 2013.

Department of Mathematics, Claremont McKenna College, Claremont, CA 91711

E-mail address: dkrumm@cmc.edu